Environmental Regulation: Monitoring and Enforcement

Prof. Richard Sweeney

ECON8852, Boston College

How should pollution regulation be enforced?

- First best: Perfect monitoring + Pigouvian taxation
- But perfect monitoring is often expensive (maybe impossible)
- In many settings, strict pollution limits are set statutorily. A regulator is tasked with identifying violators and bringing them into compliance.
- If the regulator's budget is limited, and polluters are rational, this becomes an optimal punishment problem (Becker 1968).
- Empirical question: How well does this work in practice?
  - Are regulators good at "targeting" inspections?
  - Do escalating penalties improve welfare?

At first glance, regulation does not appear effective in much of the developing world

- Emerging countries like India and China have very strict environmental standards on the books, by also have high levels of pollution.
- Why is that? Common explanations:
  - Lack of enforcement resources
  - Poor enforcement due to "corruption, laziness, or incompetence"
- "The value of regulatory discretion," by Duflo, Greenstone, Pande and Ryan (ECMA, 2018) study this in India.

Setting: Gujarat, India

- Industrialized state with severe pollution problems (all seven cities exceed national PM standards)
- Has strict standards on the books, and actually issues strict fines (9% of sample firms closed in the past)
- ... but chance of being caught still low, due to the regulator's limited inspection budget

Setting: Gujarat, India

- Industrialized state with severe pollution problems (all seven cities exceed national PM standards)
- Has strict standards on the books, and actually issues strict fines (9% of sample firms closed in the past)
- ... but chance of being caught still low, due to the regulator's limited inspection budget
- Inspection has a lot of "discretion"
  - 50% of plants inspected less than than the statutorily inspected rate
  - other plants inspected many many times more
- Is this discretion a good thing?

Setting: Gujarat, India

- Industrialized state with severe pollution problems (all seven cities exceed national PM standards)
- Has strict standards on the books, and actually issues strict fines (9% of sample firms closed in the past)
- ... but chance of being caught still low, due to the regulator's limited inspection budget
- Inspection has a lot of "discretion"
  - 50% of plants inspected less than than the statutorily inspected rate
  - other plants inspected many many times more
- Is this discretion a good thing?
  - Yes, if regulatory is using local info to target polluters
  - No, if polluters are simply bribing the inspectors

DGPR set out to evaluate discretion using an RCT

- 960 industrial plants over 2 years
- half the plants randomized to an inspection treatment
    - That was crossed with an audit treatment (Duflo et al QJE 2013)
- Treatment provided the regulator with the resources to meet the *de jure* inspection rate for all plants.
- ... It also removed any discretion in assigning additional inspections, allocating them randomly across plants
- Treatment only altered inspections, which determines if a plant starts a regulatory process which involves review, remediation and fines

The results were ... surprising

- The treatment significantly increased the inspections (2X as many as control group)

The results were ... surprising

- The treatment significantly increased the inspections (2X as many as control group)
- This in turn doubled the number of identified violations / citations.

The results were ... surprising

- The treatment significantly increased the inspections (2X as many as control group)
- This in turn doubled the number of identified violations / citations.
- However it did not cause any increase in the probability a plant was subsequently penalized.

The results were ... surprising

- The treatment significantly increased the inspections (2X as many as control group)
- This in turn doubled the number of identified violations / citations.
- However it did not cause any increase in the probability a plant was subsequently penalized.
- And did not lead to any measurable change in pollution.

END-LINE POLLUTION AND COMPLIANCE ON TREATMENTS[a]

	(1)	(2)	(3)	(4)
	Panel A. Plant-Level Costs			
	Capital Costs		Maintenance Costs	
	(USD × 10^3)	Any (=1)	(USD × 10^3)	Any (=1)
Inspection treatment (=1)	−0.221	0.0213	0.838*	0.00974
	(0.453)	(0.0344)	(0.499)	(0.0224)
Plant characteristics	Yes	Yes	Yes	Yes
Audit experiment	Yes	Yes	Yes	Yes
Control mean	2.050	0.567	0.264	0.108
Observations	791	791	791	791
	Panel B. Plant-by-Pollutant Level Pollution			
	Pollution		Compliance	
Inspection treatment (=1)	−0.105		0.0366*	
	(0.0839)		(0.0213)	
Audit treatment (=1)	−0.187**		0.0288	
	(0.0849)		(0.0258)	
Audit × inspection treatment (=1)	0.286**		−0.0365	
	(0.142)		(0.0353)	
Control mean	0.682		0.614	
Observations	4168		4168	

[a]The table shows intent-to-treat effects of inspection treatment assignment on plant costs and pollution outcomes. Panel A shows regressions for plant costs estimated at the plant level. Costs are divided into capital and maintenance costs based on descriptions of each expenditure (see Appendix A). Cost amounts are in thousands of U.S. dollars (USD). Capital costs, which are reported as lump sum in the survey, are amortized to an equivalent constant annual expenditure (using an interest rate of 20% and a 10-year equipment lifespan). Plant characteristic controls include dummies for size, use of coal or lignite as fuel, high waste water generated, and all

Why didn't the treatment reduce pollution?

Environmental Regulation: Monitoring and Enforcement

Prof. Richard Sweeney

Why didn't the treatment reduce pollution?

Treatment found many more violations, but very few *severe* violations

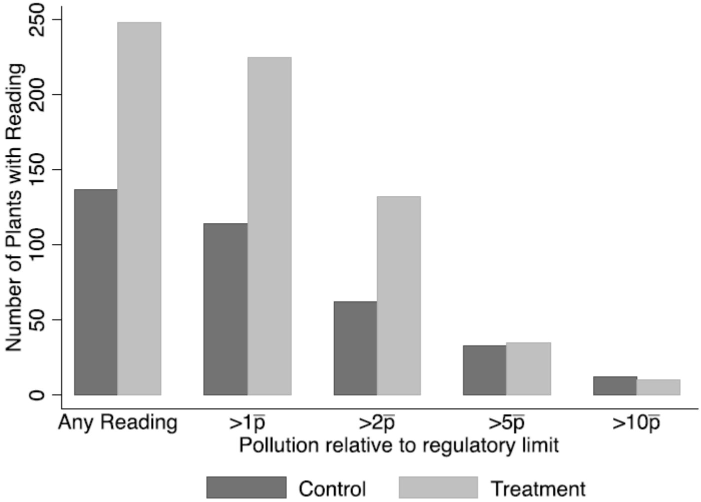

As a result, all action was very marginal

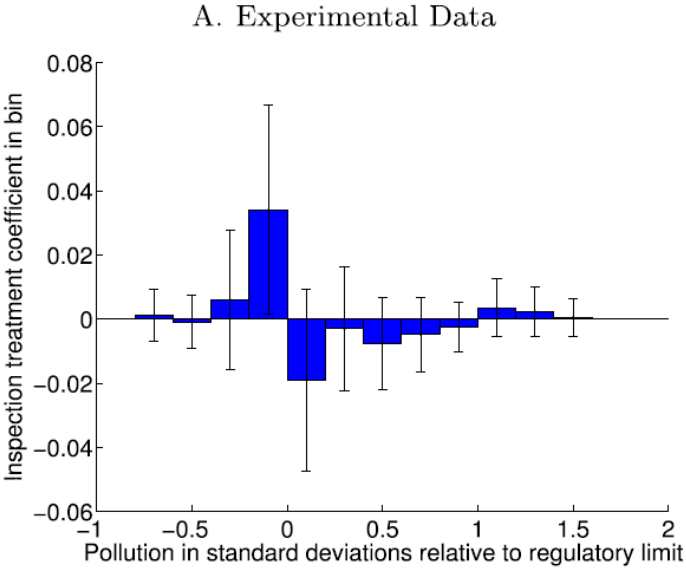

What can we say about the value of discretion from these results?

- Treatment really did two distinct things: 1) doubled inspections and 2) removed discretion.
  - Ideally they would have had two separate treatments: A) double inspections *but keep discretion*, and B) remove discretion *but keep the budget fixed* (plus maybe interaction)
  - However authors state they could not get the regulator's buy-in for B (p 2129)
  - Lacked budget to simply double discretion arm as well, and opted for interaction.

What can we say about the value of discretion from these results?

- Treatment really did two distinct things: 1) doubled inspections and 2) removed discretion.
- Zeros suggest that those two things somehow exactly offset each other? That is a bit unsatisfying
- To better understand what is really going on here, and make more precise statements about the value of discretion, DGPR estimate a structural model of regulator discretion and polluter response

Two stage game between regulator and plants

Stage 1. *Targeting*

(i) The regulator chooses an inspection targeting rule to minimize plant pollution subject to a budget of inspections.

(ii) Plants choose whether to run their abatement equipment, given their abatement cost, known level of pollution, and the regulator's targeting and penalty rules.

(iii) The regulator observes a part of plant pollution and inspects plants by applying the targeting rule (i) to this signal, yielding a pollution reading from the inspection.

Stage 2. *Penalty*

(i) The regulator acts as a *regulatory machine*, following exogenous rules for followup and punishment based on pollution measured in inspections and plant actions.

(ii) Plants face a single-agent dynamic problem: they play against the regulatory machine and decide when to comply versus when to risk future penalties.

How does firm effort relate to the level and precision of inspections?

- If firms are rational, they will abate until the marginal cost of pollution control equals the expected marginal cost of fines and remediation.
- If we knew the cost of running pollution equipment and the impact of that on pollution, we could compute this directly.
- Instead, authors infer these based on the estimated (negative) value of being in different pollution states conditional on inspection.
- Model estimated using backward induction. Delivers value function $V(p)$ for any inspector detected level of pollution level p.

STRUCTURE OF PENALTY STAGE ACTIONS[a]

	Regulatory Action				Plant Action			
Round	Inspect (1)	Warn (2)	Punish (3)	Accept (4)	Ignore (5)	Comply (6)	*N* (7)	% Left (8)
1	100.0	0.0	0.0	0.0			7423	100.0
2					99.6	0.4	7423	
3	1.0	9.5	2.2	87.3			7423	100.0
4					92.8	7.2	941	
5	23.3	4.8	5.3	66.6			941	12.7
6					91.1	8.9	314	
7	18.8	11.8	9.9	59.6			314	4.2
8					83.5	16.5	127	
9	21.3	5.5	18.1	55.1			127	1.7
10					82.5	17.5	57	
11	26.3	3.5	10.5	59.6			57	0.8
12					87.0	13.0	23	
13	26.1	4.3	8.7	60.9			23	0.3
14					77.8	22.2	9	
15+	16.7	8.3	0.0	75.0	100.0	0.0	9	0.1
Total without inspections	0.0	4.6	1.6	42.7	50.2	0.9	7824	
Total	31.0	3.2	1.1	29.4	34.6	0.6	25,217	

[a]The table reports actions taken by the regulatory machine and by the firm in the penalty stage using administrative data. Figure 5 defines actions and their payoffs and Table II maps them to regulatory documents. Each of columns 1–6 gives the probability, within that row, of the party moving at that round when taking the action indicated in the column header. Column 7 gives the total number of actions observed in that round and column 8 gives the percentage of penalty stages that continue up to at least that round. The

Targeting stage model

- Firms have latent pollution:
  $\ln \tilde{P}_j m = \phi_0 + \phi_1 X_j + \mu_{1j} + \mu_{2jm}$
  where μ_1 is common info and μ_2 is known only by the plant.
- Can run pollution equipment at cost c. This reduces pollution proportionally
  $\ln P_{jm} = \ln \tilde{P}_j m + \phi_2 Run$
- Run if: $I_j(V_0(P_{jm}) - V_0(\tilde{P}_j m)) > c_j$
- DGPR model this as a scaled probit.
  - "inner" error reflects common info μ_1
  - outer param scales level of total inspections up and down.

Estimates of Targeting Stage Parameters[a]

	Constrained		Unconstrained	
	Initial Inspections (1)	Log Pollution (2)	Initial Inspections (3)	Log Pollution (4)
Panel A. Targeting and Pollution Equations				
Inspection treatment	0.095		0.162	
	(0.009)		(0.025)	
Run equipment (=1)		−1.902		−0.711
		(0.160)		(0.308)
Inspection targeting shift parameter (λ_1)	−0.395		−0.220	
	(0.003)		(0.066)	
Inspection targeting level parameter (λ_2)	33.022		10.064	
	(1.876)		(3.137)	
Constant		0.212		−0.009
		(0.109)		(0.102)
Panel B. Distributions of Pollution and Maintenance Cost Shocks				
Standard deviation of observed pollution shock (σ_1)	0.069		0.111	
	(0.003)		(0.022)	
Standard deviation of unobserved pollution shock (σ_2)	1.033		0.864	
	(0.047)		(0.042)	
Mean of log maintenance cost (μ_c)	2.388		1.833	
	(0.061)		(0.334)	
Panel C. Test of Targeting Optimality Constraints				
Distance metric test statistic χ^2_2	16.1039			
Test p-value	0.0003			

[a]The table reports parameters of the targeting stage of the model. The first two columns 1 and 2 report estimates from the constrained model where the regulator is constrained to target optimally based on observed pollution shocks, and columns 3 and 4

What did the experiment buy them?

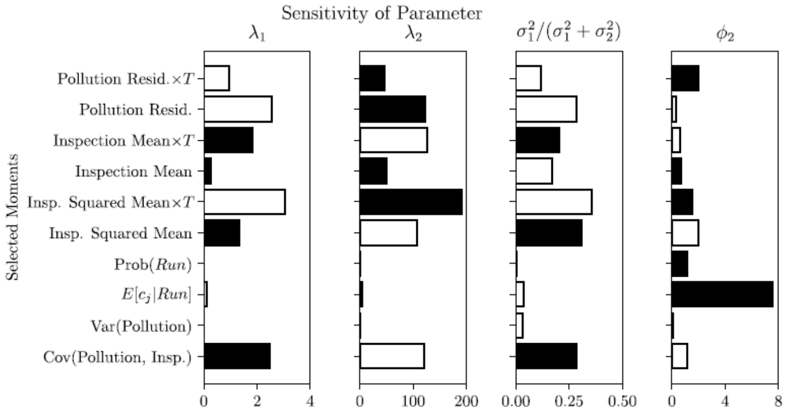

FIGURE 6.—Sensitivity of targeting parameters to moments. The figure shows the Andrews, Gentzkow, and Shapiro (2017) sensitivity measure of selected targeting model parameter estimates with respect to selected moments used to estimate the model. The panels, left to right, show the sensitivity of the parameters or functions of parameters λ_1, λ_2, and $\sigma_1^2/(\sigma_1^2+\sigma_2^2)$, with respect to moments indicated by the row headers. The length of each bar is the local sensitivity of the parameter with respect to the row moment, measured in units of the parameter value per 1 standard deviation change in the moment. From the rows, we omit the products of pollution residuals and inspection means with observable plant characteristics other than treatment status.

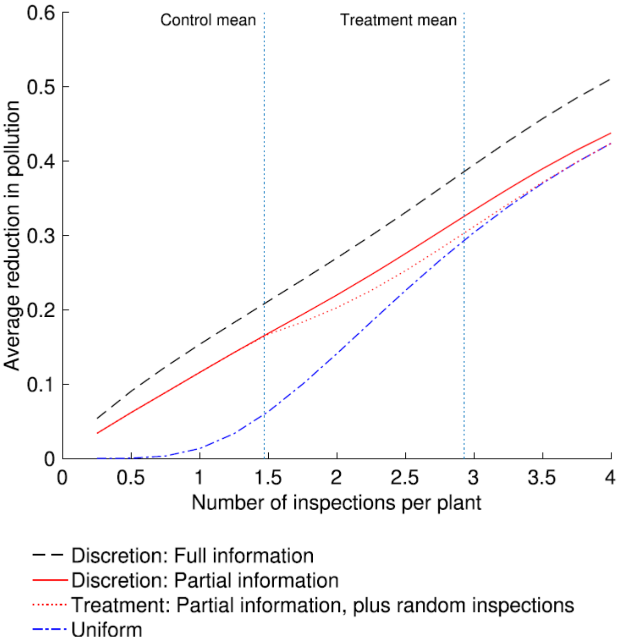
0.6
0.5
0.4
0.3
0.2
0.1
0
Average reduction in pollution
Control mean
Treatment mean
0
0.5
1
1.5
2
2.5
3
3.5
4
Number of inspections per plant
Discretion: Full information
Discretion: Partial information
Treatment: Partial information, plus random inspections
Uniform

Duflo wrapup

- Regulation in India does not seem very effective, and there are competing ideas about why.
- Setup an RCT to study, and found a surprising result, but implications for regulation were unclear.
- Use structure + experimental variation to estimate a very rich model of regulator discretion and polluter response.
- Find that, in this setting, discretion is actually quite helpful.
- Real problem is that the regulator has a very limited budget, and observed only a very weak signal about plant pollution

Dynamic Regulation

- In India fines were actually quite high, but not high enough to offset the very low inspection probability. As a result compliance was low.
- In the US, Harrington (JPubEcon, 1988) noted a "paradox"
  - Fines are quite low, with the expected penalty much lower than the expected remediation cost.
  - However compliance is quite high.
- Harrington shows that this is rationalized by a repeated-game model in which the regulator applies escalating scrutiny.

Why not just go the Becker (1968) route?

Harrington has a nice discussion of the limits

- Unlike crime, many pollution violations are not willful
- Severe but rarely-imposed penalties might seem capricious and unfair.
- There is an upper limit to the tine that can be imposed on any given firm such that the firm is not driven into bankruptcy

Intuition for the value of dynamic escalation

- Imagine you show up at a firm and discover a violation.
  - This may (or may not) be news to the firm too
  - That external damage is a sunk cost.
  - If the regulator dislikes fines, penalizing for this past externality now strictly decreases welfare.
- What matters once we are in this state is that the plant remedy the situation. This can be done with a promise to both increase the frequency of visits and severely punish the next violation.
- Harrington's simple model shows this escalating policy *can* achieve any pollution outcome at lower levied fines, in theory.

How valuable is dynamic escalation in practice?

- Would need to the costs of enforcement (fines, inspections)
- Need to know how escalating fines effects realized pollution, and the costs of the remediation efforts it induces.
- With those primitives, could consider other counterfactual policies, and compare welfare and pollution outcomes.
- Brings us to "Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations," by Blundell, Gowrisankaran, and Langer (AER, 2020)

Setting: Clean Air Act

Figure 1: EPA Clean Air Act Amendment Enforcement by Regulatory Status

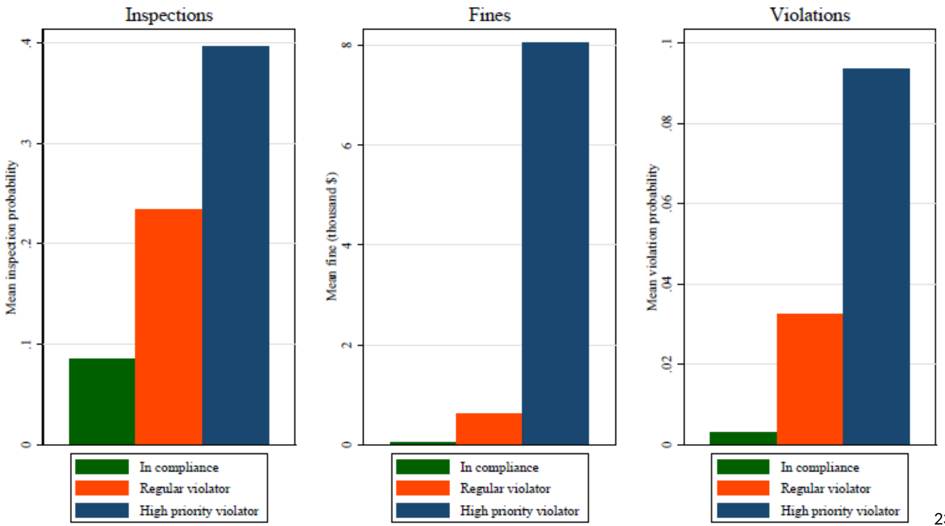

What does this look like in practice?

Nice example from a refinery in Texas

- Refinery had a history of little violation
- 2011 they left a valve open, accidentally leaking VOCs and benzene
- Due to the severity of this, became a high priority violator (HPV), instigating higher scrutiny and fines
- 2012 another minor leak happens, but this time fines were doubled.
- This continued until the firm made two large investments in abatement and monitoring equipment.
- At this point the firm's HPV designation was removed.

BGL model this as a dynamic game

- Regulator pre-commits to policy function (similar to Duflo's regulatory "machine")
- Regulator wants to maximize compliance, but inspections are costly and dislikes imposing fines.
- Violations arise stochastically, and plants detect them with the regulator.
- Plants rationally decide wether or not to remediate them with costly investment.
    - These investments are not always successful.
- Plants are heterogeneous but the regulator cannot contract on these differences

BGL model this as a dynamic game

- Regulator pre-commits to policy function (similar to Duflo's regulatory "machine")
- Regulator wants to maximize compliance, but inspections are costly and dislikes imposing fines.
- Violations arise stochastically, and plants detect them with the regulator.
- Plants rationally decide wether or not to remediate them with costly investment.
    - These investments are not always successful.
- Plants are heterogeneous but the regulator cannot contract on these differences

What data would you like to estimate a model like this?
What data do BGL actually have?

Primary data: EPA ECHO Database

- Plant info including location
- Quarterly info on enforcement actions
- Violations become "resolved" when ECHO records a *Prevention of Significant Deterioration* (PSD). This is all they see regarding firm investment.

Primary data: EPA ECHO Database

- Plant info including location
- Quarterly info on enforcement actions
- Violations become "resolved" when ECHO records a *Prevention of Significant Deterioration* (PSD). This is all they see regarding firm investment.

Facts from this data

- 2,355,908 plant/quarter observations.
- 95.6 percent of observations indicate compliance.
- Investment occurs in 4.9% of violator quarters, 17.5% HPV.
- 88.4% of *plants* are never out of compliance
- Only 4.2% of plants ever HPV

BGL Markov Model

- Ω_t contains payoff relevant state vars
  - past (discounted) violations, status, investment, region, sector, gravity (of damages)
- Inspection probability $\mathcal{I}(\Omega)$
  - $Ins(\Omega_t)$ denotes actual inspection decision
- e_t is a signal about plant pollution conditional on inspection
- These then feed into
  - probability of new violation $Vio(\Omega, e)$
  - punishment $Fine(\Omega, e)$
  - transition function $T(\Omega, e) = \tilde{\Omega}$
- Firm decides whether to take costly investment action $X \in (0, 1)$ conditional on learning they are not in compliance ($Com(\Omega, e) = 0$)

Taking this to data: Plants

Plant flow utility from regulatory actions:

$$U(\Omega, e) = \theta^I Ins(\Omega) + \theta^F Fine(\Omega, e) + \theta^V Viol(\Omega, e) + \theta^H HPV(T(\Omega, e))$$

If found in violation, firms can also make costly investments: $\theta^X + \epsilon_{xt}$

Structural parameters: $\Theta = (\theta^I, \theta^F, \theta^W, \theta^H, \theta^X)$
– Authors expect these to vary across firm, but assume that this heterogeneity is not something the regulator can contract on

Taking this to data: Plants

Plant flow utility from regulatory actions:

$$U(\Omega, e) = \theta^I Ins(\Omega) + \theta^F Fine(\Omega, e) + \theta^V Viol(\Omega, e) + \theta^H HPV(T(\Omega, e))$$

If found in violation, firms can also make costly investments: $\theta^X + \epsilon_{xt}$

Structural parameters: $\Theta = (\theta^I, \theta^F, \theta^W, \theta^H, \theta^X)$
– Authors expect these to vary across firm, but assume that this heterogeneity is not something the regulator can contract on

Note: Many of these things observed in DGPR.

Taking this to data: Regulator

- In order to make rational decisions, plants need expectations of regulator actions
- BGL model these as CCP's
    - Inspection probability $\mathcal{I}(\Omega)$ (probit, 2 states)
    - Probability of new violation $Vio(\Omega, e)$ (probit, 2 states)
    - Punishment $Fine(\Omega, e)$ (Tobit, 20 values)
    - Transition violator status $HPV(\Omega, e)$ (mlogit, 3 states)
- Gives them probabilities of all 240 possible outcomes from every state.
- Note: signal e is now the structural error in each of these equations.

Can now compute dynamic expectations

Value function at the beginning of each period (*before* any action)

$$V(\Omega) = \sum_{I \in (0,1)} \mathcal{I}(\Omega)^{I}(1 - \mathcal{I}(\Omega))^{1-I} \int [U(\Omega, e) + \tilde{V}(T(\Omega, e))] dP(e|Ins, \Omega)$$

Can now compute dynamic expectations

Value function at the beginning of each period (*before* any action)

$$V(\Omega) = \sum_{I \in (0,1)} \mathcal{I}(\Omega)^I (1 - \mathcal{I}(\Omega))^{1-I} \int [U(\Omega, e) + \tilde{V}(T(\Omega, e))] dP(e | Ins, \Omega)$$

Let $\tilde{V}(\tilde{\Omega})$ denote the value function at the point right after the regulator has moved but before the plant receives its draws of e.

Conditional on inspection, continuation value is

$$\tilde{V}(\tilde{\Omega}) = Com(\tilde{\Omega}) \times \int [\beta V(\tilde{\Omega}, \theta) + \epsilon_0] dF(\epsilon_0) + (1 - Com(\tilde{\Omega}))$$

$$\times \int \int \max[\beta V(\Omega | \tilde{\Omega}, X = 0) + \epsilon_0, -\theta^X + \beta V(\Omega | \tilde{\Omega}, X = 1) + \epsilon_1\}$$

Estimation: Homogeneous case

1. Guess θ
2. Use functions $V(\Omega)$, $\tilde{V}(\tilde{\Omega})$, and CCPs to get value function (Rust (1987))
3. search over θ to max quasi-log-likelihood $\log L(\theta) =$

$$\sum_i \sum_t \log\left([X_{it} Pr(X=1)|\tilde{\Omega}_{it},\theta) + (1-X_{it})(1-Pr(X=1)|\tilde{\Omega}_{it},\theta)]\right)$$

where $Pr(X=1|\tilde{\Omega},\theta)$ comes from the investment probabilities at the fixed point of the Bellman

4. SEs bootstrapped

Estimation: Random coefficient

- θ could vary across plants
- Typical approach (ie BLP)
  - Specify a parametric distribution $F(\theta)$
  - Guess μ, σ, take draws, compute likelihood of data.
- What do BGL do instead?

Estimation: Random coefficient

- Rather than specifying this distributions $F(\theta)$, BGL implement a non-parametric approach from Fox, Kim, Ryan and Bajari (QE)
  1. Specify many θ_j plant types.
  2. Assign each type a probability η_j
  3. Generate moments based on the likelihood that a plant is a given type, and the optimal decision given that type.
- Why do this?

Estimation: Random coefficient

- Rather than specifying this distributions $F(\theta)$, BGL implement a non-parametric approach from Fox, Kim, Ryan and Bajari (QE)
  1. Specify many θ_j plant types.
  2. Assign each type a probability η_j
  3. Generate moments based on the likelihood that a plant is a given type, and the optimal decision given that type.
- Why do this?
- Main computational burden comes from solving $V(\theta)$ and $\tilde{V}(\theta)$.
- Rather than doing this in an innerloop for every guess at the hyperparameters of the $F_\gamma(\theta)$ distribution, then can do this ONCE at the beginning
- Why don't we do this more often?

Estimation Steps : Random coefficient

1. Set θ state space (they pick 10,000 grid points)
2. Solve the value function for each θ_j
3. Calculate three moments $m_k(\theta_j)$
    1. Probability of observed states
    2. Probability of observed state - investment pairs
    3. Probability of a 6 period sequence of plant investments conditional on type
4. M1 and M2 require Assumption 2: observed data on compliance related state vars (ω_2) conditional on other state vars (ω_1).
5. Minimize distance between observed moments in the data and weighted sum $\sum_j \eta_j m(\theta_j)$
6. Standard errors **on counterfactuals** constructed using bootstrap

Estimated parameters

Table 3—Estimates of Plants' Structural Parameters

	Quasi-likelihood estimates	GMM random coefficient estimates (1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	−2.872 (0.041)	−2.334	−1.326	−2.498	−2.540	−1.988	0.153
Inspection utility (θ^I)	−0.049 (0.049)	−0.194	0.444	−0.096	0.897	0.001	−2.483
Violation utility (θ^V)	−0.077 (0.197)	0.143	0.128	0.650	−0.100	−2.169	−2.006
Fine utility (millions \$, θ^F)	−5.980 (1.005)	−5.181	−6.073	−6.766	−8.460	−7.494	−7.524
HPV status utility (θ^H)	−0.065 (0.015)	−0.029	−0.234	−0.078	−0.411	0.070	−2.437
Weight on parameter vector	1	0.438	0.174	0.170	0.126	0.049	0.019

Notes: For the quasi-likelihood approach, we estimate the costs themselves, whereas for the GMM random coefficient approach, we estimate the weights (in the bottom row) on each potential vector of costs. For GMM estimates, we report the six parameter vectors with the highest weight. Standard errors for quasi-likelihood estimates, which are bootstrapped with resampling at the plant level, are in parentheses.

Counterfactual Preliminaries

- In the model, actions depend observed states Ω_t **and the environmental compliance signal** e_t
- Problem, e_t not observed.
- **Assumption 1:** Signal (e) is a function of the regulatory state Ω, regulatory machine inspection CCP's ($\mathcal{I}()$), and the inspection decision (Ins_t).
- Note that is does NOT depend on the fine schedule.
- Authors state this is stronger than they need *for estimation*. Alternative, could have given firm's priors about $F(e|\omega)$
- But A1 is critical for counterfactuals

Counterfactual Preliminaries

- Signals must be at least partially related to actions firms take. Thus, their distribution depends on the regulatory program (as in Duflo).
- In order to compute optimal behavior, BGL need the CCPs. So they must assume that those are the same in any counterfactual (remember that e is the structural error in those).
- But, A1 says that e does not depend on fines. It implies that two plants with the same history of violations facing the same inspection probabilities will face the same distribution of e, *even if they face different fine schedules*.
- This allows them to change the fine schedule in counterfactuals, but not the frequency of inspections.

What do people think about Assumption 1, given the previous paper?

What do people think about Assumption 1, given the previous paper?

Presumably if escalating penalities improved probability of compliance, and inspections are costly, the optimal policy somewhat reduces the latter.

– so this slightly underestimates the gains from dynamic enforcement (I think, but I wish there was more discussion of this in the paper.)

Counterfactuals - part 1

Initial motivation: How effective is dynamic escalation?

1. Keep baseline total fines the same, but don't escalate dynamically (static fines)
2. same thing but hold POLLUTION constant – how much higher do fines need to be?
3. double HPV fines

Table 4—Counterfactual Results: Changing the Escalation Rate of Fines

	Data (1)	Baseline (2)	Same fines for all violators; fines constant (3)	Same fines for all violators; pollution damages constant (4)	Fines for HPVs doubled relative to baseline (5)
Compliance (percent)	95.62	95.11 (0.22)	66.72 (13.91)	94.49 (0.62)	95.52 (0.24)
Regular violator (percent)	2.88	3.47 (0.25)	2.53 (0.57)	2.72 (0.56)	3.47 (0.26)
HPV (percent)	1.50	1.42 (0.05)	30.75 (14.43)	2.79 (0.65)	1.01 (0.03)
Investment rate (percent)	0.40	0.54 (0.05)	0.47 (0.06)	0.65 (0.09)	0.55 (0.05)
Inspection rate (percent)	9.65	9.41 (0.06)	20.54 (5.41)	9.88 (0.23)	9.28 (0.05)
Fines (thousands $)	0.18	0.32 (0.03)	0.32 (0.03)	1.98 (1.62)	0.36 (0.03)
Violations (percent)	0.55	0.54 (0.01)	5.00 (2.20)	0.74 (0.10)	0.49 (0.01)
Plant utility	—	0.006 (0.034)	0.077 (0.091)	0.001 (0.039)	0.005 (0.034)
Pollution damages (milions $)	1.65	1.53 (0.03)	4.04 (1.19)	1.53 (0.03)	1.48 (0.02)

Notes: Each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column 1 presents the value of each statistic in our data. Column 2 presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines and the HPV cost faced by plants. Columns 3 and 4 impose the same fines for all regular and high priority violators for a given fixed state. Column 5 doubles the fines for plants in HPV status. All values are per plant/quarter. Bootstrapped standard errors are in parentheses.

Counterfactuals - part 2

Counterfactuals

1. Pigouvian fines based on damages (no escalation)
2. 3X pigouvian damages (did not understand this one)
3. Fines set to keep damages constant

TABLE 5—COUNTERFACTUAL RESULTS: SCALED PIGOUVIAN FINES

	Baseline (1)	Pigouvian fines (2)	Pigouvian fines scaled by 1/3 (3)	Pigouvian fines scaled to yield base pollution damages (4)
Compliance (percent)	95.11 (0.22)	96.69 (1.05)	95.38 (1.78)	82.44 (4.60)
Regular violator (percent)	3.47 (0.25)	1.60 (0.30)	2.09 (0.30)	2.88 (0.37)
HPV (percent)	1.42 (0.05)	1.72 (1.02)	2.52 (1.80)	14.68 (4.89)
Investment rate (percent)	0.54 (0.05)	0.86 (0.05)	0.79 (0.06)	0.53 (0.06)
Inspection rate (percent)	9.41 (0.06)	9.34 (0.33)	9.60 (0.58)	14.18 (1.72)
Fines (thousands $)	0.32 (0.03)	55.24 (1.81)	19.06 (0.69)	1.58 (1.67)
Violations (percent)	0.54 (0.01)	0.52 (0.12)	0.60 (0.21)	2.31 (0.60)
Plant utility	0.006 (0.034)	−0.349 (0.047)	−0.117 (0.038)	0.032 (0.042)
Pollution damages (millions $)	1.53 (0.03)	1.32 (0.02)	1.32 (0.02)	1.53 (0.03)

Notes: Each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column 1 presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines faced by plants. All values are per plant/quarter. Bootstrapped standard errors are in parentheses.

Discusion

What did and didn't you like about this paper?

What questions left open?

Other applications?